

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
APPLICATION FOR LETTERS PATENT

Linked Value Replication

Inventor(s):

William B. Lees
Jeff B. Parham
Mark R. Brown
Donald J. Hacherl

ATTORNEY'S DOCKET NO. MS1-677US

1 **RELATED APPLICATION**

2 This application claims priority to U.S. Provisional Application No.
3 60/212950, filed June 21, 2000, entitled "Link Value Replication", to Brown et al.

5 **TECHNICAL FIELD**

6 This invention relates to network systems and, in particular, to linked multi-
7 valued object attribute replication in a network-wide directory service.

9 **BACKGROUND**

10 In a network-wide directory service maintaining objects having multi-
11 valued attribute lists, such as a mail distribution list or a personnel list for a
12 security-based system, simultaneous updates from more than one networked data-
13 entry site can cause a replication conflict. For example, Active Directory™ is an
14 enterprise-wide directory service in Windows® 2000 using a state-based, multi-
15 master replication model that is susceptible to replication conflicts with respect to
16 its object store structure. Windows® 2000 is an operating system licensed by
17 Microsoft Corporation of Redmond, Washington.

18 In a network-wide partitioned directory, each domain controller in a
19 separate domain of the network maintains a copy of a partition of the directory
20 which typically contains those objects that are pertinent to only a particular
21 domain. Replication defines that a change to a directory made on one computer
22 will change the directory on all computers in a network having a replica of the
23 directory. A copy of the contents of one directory partition on a specific domain
24 controller is identified as a replica. Replication updates replicas among the
25 domain controllers that store the same directory partitions. Convergence defines

1 that if a network system is allowed to reach a steady state in which no new updates
2 are occurring, and all previous updates have been completely replicated, all
3 replicas ideally converge to the same set of values.

4 A multi-master replication model defines that several servers (e.g., the
5 domain controllers) in a network system can contain writeable replicas of an
6 object that is intended to be kept consistent between the servers. Master replicas
7 accept updates independently without communicating with other master replicas.
8 If updates cease and replication continues, all replicas of an object at each server
9 will ideally be updated to the same value. Replication propagates changes made
10 on any specific domain controller to all other domain controllers in the network
11 that store the directory partition in which a change occurs.

12 A state-based replication model defines that each master applies updates,
13 both originating and replicated, to its replica as they arrive. Replication is derived
14 from the current state of the source replica at hand. Each directory partition
15 replica stores per-object and per-attribute data to support replication.

16 An alternative to a state-based replication model is a log-based replication
17 model. In a conventional log-based replication system, each master server keeps a
18 log of any updates that it originates. When replicating, each master server
19 communicates its log to every other replica. When receiving a log at a replica, the
20 replica applies the log, bringing its own state more up-to-date.

21 With a conventional state-based replication model, there can be conflicts
22 with object attribute value updates because the lowest level of granularity for
23 updates is at the attribute level of an object, and not at the attribute value level.
24 Even though an attribute may contain multiple values (i.e., a multi-valued
25 attribute), all of the values are considered as a single unit for the purpose of

replication. The following example, described with reference to Figs. 1 and 2, illustrates the occurrence of a replication conflict when implementing a network-wide directory service with a conventional state-based replication model.

Fig. 1 shows a network architecture 100 having a directory service that maintains objects associated with a mail distribution list. The network 100 has a first domain controller 102, computer A, and a second domain controller 104, computer B, that are interconnected via a communications network 106. Computer 102 has a directory 108 that stores a mail group 110(A) which has multiple associated group objects, such as object 112(A). Group object 112(A), identified as object M, is associated with mail group 110(A) and identifies the individual recipients of a mail distribution list in the mail group.

Computer 104 has a directory 114 which is a replica of directory 108 in computer 102. Directory 114 stores a mail group 110(B) which has an associated group object 112(B), also identified as object M because it is a replica of object 112(A) stored in directory 108 at computer 102.

The group object 112 has a data structure 116 that illustrates data stored in the object. The data structure 116 stores object properties, identified as attributes 118, and attribute values for each attribute, identified as metadata 120. The object 112 has a name attribute 122 that identifies an association with mail group 110. Metadata 124 indicates the association with the mail group and also includes a latest version number and an update timestamp for the name attribute 122. The version number, v1, indicates a first version of the name attribute 122 and the timestamp, t1, indicates when the first version of the attribute was created.

The object 112 has an identifier attribute 126 that associates a global unique identifier (GUID) in metadata 128 for the object. Each instance of the object,

1 112(A) and 112(B), has a different and unique GUID within network 100.
2 Metadata 128 also includes a latest version number, v1, and an update timestamp,
3 t1, for the identifier attribute 126.

4 The object 112 also has a multi-valued members attribute 130 that
5 associates the individual recipients in the mail distribution list. Metadata 132 for
6 the members attribute includes a latest version number, v1, and an update
7 timestamp, t1. Metadata 132 also includes a link table reference to a data structure
8 134. Link table 134 maintains the linked values (e.g., the recipients in the mail
9 distribution list) for the multi-valued members attribute 130.

10 Link table 134 identifies the object owning the link table at source 136
11 which indicates that object M owns the link table. Each recipient in the mail
12 distribution list is identified as a referenced object at destination 138 which, in this
13 example, indicates two recipients. Link table 134 also identifies the associated
14 object attribute for each destination 138 at linkID 140. In this example, linkID
15 140 identifies that each recipient 138 is associated with the members attribute 130.

16 If the list of recipients 138 is changed on computer A, then computer B
17 needs to be updated with the changes. During replication, computer A sends
18 computer B the entire contents of the members attribute 130, which includes the
19 entire link table 134, because the lowest level of granularity for conventional
20 replication updates is at the attribute level of an object, and not at the attribute
21 value level. Although only a single value within the members attribute value list
22 may be changed (i.e., a recipient is deleted, added, and/or updated), computer A
23 cannot convey to computer B which recipient has changed. Computer A can only
24 convey that some value in the members attribute 130 has been changed.

1 The problem is compounded for a large number of attribute values and by
2 the scale of the network. Computer B can only receive the entire contents of the
3 members attribute 130 and either compare the new object attribute with what
4 computer B has stored locally to update the change, or computer B can delete its
5 entire local copy of the members attribute and update the attribute with the new
6 copy of members from computer A. Either case presents an efficiency problem
7 for computer B. The problem is further compounded for multiple networked sites
8 each having replica to be updated.

9 Furthermore, a conflict occurs during replication when a multi-valued
10 object attribute, such as members, is updated at different networked sites within a
11 relatively short amount of time before a scheduled replication. This is identified
12 as a replication latency period. Changes made to a multi-valued attribute
13 simultaneously, or within the replication latency period, can cause a replication
14 convergence conflict that will result in the loss of a data update.

15 If two independent attribute changes converge from different networked
16 sites, and a first attribute change prevails in a conflict resolution over a second
17 attribute change, then the values of the first attribute change will replace all of the
18 values of the second attribute change. This policy is acceptable for an attribute
19 that is single-valued, or when it makes sense to change all of the values of an
20 attribute together as a group. However, replication conflicts can result in lost data
21 when it is desirable that individual values of a multi-valued object attribute
22 replicate independently.

23 Fig. 2. continues the example and illustrates how a replication conflict can
24 occur between two objects having updated multi-valued attributes and how
25 resolution of the conflict can result in the loss of one of the data updates. Initially,

1 as shown in Fig. 1, computer A has an object 112(A) with a multi-valued members
2 attribute 130. The attribute has two values, recipient1 and recipient2, in link table
3 134. Computer B also has an up-to-date replica of object M.

4 In Fig. 2, a data administrator at computer A deletes recipient1 from the
5 mail distribution list 138(A) in link table 134(A) and, as illustrated, recipient1 no
6 longer exists. The data administrator also adds a new recipient3 to the mail
7 distribution list 138(A) as indicated by 200. Metadata 132(A) for members
8 attribute 130(A) is updated to version2 (v2) of the mail distribution list occurring
9 at time2 (t2) as indicated by 202.

10 Within a replication latency period, such as five minutes or less, for
11 example, a second data administrator at computer B adds a new recipient4 to the
12 mail distribution list 138(B) as indicated by 204. Metadata 132(B) for members
13 attribute 130(B) is updated to version2 (v2) of the mail distribution list occurring
14 at time3 (t3) as indicated by 206.

15 When computers A and B replicate directories 108 and 114, respectively,
16 there will be a replication conflict because the members attribute was updated at
17 both network sites during a replication latency period. Conventionally, the
18 conflict can be resolved by a policy that allows the most frequent writer to prevail
19 first followed by the last writer prevails. That is, the higher version number
20 prevails first, followed by the latest timestamp. In the example, both network sites
21 have a version2 (v2) in metadata 132 for members attribute 130. Thus, computer
22 B wins the replication conflict because the latest timestamp is time3 (t3) which is
23 later than time2 (t2) at computer A. Other resolution policies may resolve
24 replication conflicts with only a version number, or with only a timestamp.

1 To replicate, computer A updates metadata 132(A) for members attribute
2 130(A) by replacing all of the values for the attribute. That is, the entire link table
3 134(A) is replaced in directory 108 in computer A with link table 134(B) from
4 computer B. Although not shown specifically, the resultant replica for object 112
5 at both of the network sites is that shown for computer B. The mail distribution
6 list at both computers A and B (i.e., the recipient values 138) will include
7 recipient1, recipient2, and recipient4. The update at computer A to remove
8 recipient1 and add recipient3 is lost in the resolution of the replication conflict.

9 Simultaneous attribute updates at different networked sites can cause a
10 replication convergence that requires a conflict resolution in a state-based
11 replication model because objects are not necessarily replicated in the order in
12 which they are updated. Replication conflicts arise because the lowest level of
13 granularity for updates is at the attribute level of an object, and not at the attribute
14 value level. Even though an attribute may contain multiple values, all of the
15 values are considered as a single unit for the purpose of replication. Updates to
16 individual values of multi-valued attributes need to be accounted for during
17 replication to avoid a replication conflict that results in lost data.

18

19 **SUMMARY**

20 A network system domain controller maintains a directory of objects
21 having multi-valued attributes. The attributes have multiple linked values and the
22 individual values have conflict-resolution data that indicates a change to an object
23 at an attribute-value level. The conflict-resolution data includes a version number
24 that identifies a latest version of an individual value, an update timestamp that

1 identifies when an individual value is updated or changed, and a creation
2 timestamp that identifies when an individual value is created.

3 A second network domain controller stores a replica of the directory in
4 which a replica of the objects is maintained. The domain controllers replicate the
5 objects in the directories and update the individual linked values of the attributes.
6 Replication conflicts are identified and resolved with the conflict-resolution data at
7 the attribute-value level of the objects. Additionally, the individual values have an
8 associated deletion timestamp that either indicates the existence of a value in an
9 object, or indicates that a particular value has been identified to be deleted from a
10 multi-valued attribute.

11

12 **BRIEF DESCRIPTION OF THE DRAWINGS**

13 The same numbers are used throughout the drawings to reference like
14 features and components.

15 Fig. 1 illustrates an example of conventional state-based replication.

16 Fig. 2 illustrates an example of conventional state-based replication.

17 Fig. 3 is a block diagram of a network architecture.

18 Fig. 4 illustrates data structures in the Fig. 3 network architecture.

19 Fig. 5 illustrates data structures in the Fig. 3 network architecture.

20 Fig. 6 illustrates data structures in the Fig. 3 network architecture.

21 Fig. 7 illustrates data structures in the Fig. 3 network architecture.

22 Fig. 8 illustrates data structures in the Fig. 3 network architecture.

23 Fig. 9 illustrates a network architecture and a data structure.

24 Fig. 10 illustrates data structures in the Fig. 9 network architecture.

25 Fig. 11 illustrates data structures in the Fig. 9 network architecture.

1 Fig. 12 is a flow diagram of a method for replicating multi-valued object
2 attributes.

3 Fig. 13 is a diagram of a computing system and environment that can be
4 utilized to implement the technology described herein.
5

6 **DETAILED DESCRIPTION**

7 The following technology describes systems and methods to individually
8 replicate multi-valued object attributes. A linked value replication model
9 described herein replicates attribute values individually for multi-valued object
10 attributes and reduces the possibilities of replication conflicts when the attribute
11 values converge at all replicas within a network.

12 Fig. 3 shows a network architecture 300 having any number of domain
13 controllers 302(1...n) that implement a distributed network-wide directory service
14 and that are interconnected via a communications network 304. The network
15 domain controllers 302 locally administrate the network 300 at a particular
16 network branch site. Network domain controller 302(1) is an exemplary
17 computing device of the other domain controllers (i.e., 302(2...n)) in the network
18 300. The domain controllers 302 have a processor 306 and a memory 308. The
19 memory 308 stores a directory service 310 that is executable on the processor 306.

20 The memory 308 also stores a directory 312 of any number of objects
21 314(1...x) that are distributed among the domain controllers 302. An update or
22 change to an object 314 at any one domain controller can be replicated to any of
23 the other domain controllers in the network 300 that store a copy of the same
24 object 314. The domain controllers 302 communicate replication changes via the
25 communications network 304. See the description of “Exemplary Computing

1 System and Environment" below for specific examples of the network
2 architectures and systems, computing systems, and system components described
3 herein.

4 Fig. 4 shows an example of object data structures in network architecture
5 300. Network 300 has a first domain controller A, identified as 302, and a second
6 domain controller B, identified as 316, that are interconnected via the
7 communications network 304. Domain controller A has a directory 312 that stores
8 a security group 318(A) which has multiple associated group objects, such as
9 object 314(A). The group object 314(A), identified as object S, is associated with
10 the security group 318(A) and identifies individual accounts in a security list.

11 Domain controller B has a directory 320 which is a replica of directory 312
12 in domain controller A. Directory 320 stores a security group 318(B) which has
13 an associated group object 314(B), also identified as object S because it is a
14 replica of object 314(A) stored in directory 312 at domain controller A.

15 The group object 314 has a data structure 320 that illustrates data stored in
16 the object. The data structure 320 stores object properties, identified as attributes
17 322, and attribute values for each attribute, identified as metadata 324. The object
18 314 has a name attribute 326 that identifies an association with security group 318.
19 Metadata 328 indicates the association with the security group and also includes a
20 latest version number and an update timestamp for the name attribute 326. The
21 version number, v1, indicates the first version of the name attribute 326 and the
22 timestamp, t1, indicates when the first version of the attribute was created.

23 The object 314 has an identifier attribute 330 that associates a global unique
24 identifier (GUID) in metadata 332 for the object. Each instance of the object,
25 314(A) and 314(B), has a different and unique GUID within network 300.

1 Metadata 332 also includes a latest version number, v1, and an update timestamp,
2 t1, for the identifier attribute 330.

3 The object 314 also has a multi-valued members attribute 334 that
4 associates the individual accounts in the security list. Metadata 336 for the
5 members attribute does not include a latest version number and update timestamp
6 for reasons that will become apparent below. Metadata 336 includes a link table
7 reference to a data structure 338. Link table 338 maintains the linked values (e.g.,
8 the accounts in the security list) for the multi-valued members attribute 334.

9 Link table 338 identifies the object owning the link table at source 340
10 which indicates that object S owns the link table. Each account in the security
11 personnel list is identified as a referenced object at destination 342 which, in this
12 example, indicates two accounts. Link table 338 also identifies the associated
13 object attribute for each destination 342 at linkID 344. In this example, linkID
14 344 identifies that each account 342 is associated with the members attribute 334.

15 The linked values (i.e., accounts 342) of the members attribute 334 are like
16 virtual attributes in that the values have identifying and defining data and exist in
17 the context of the containing object. Link table 338 maintains valuedata 346 for
18 each account 342 that includes a latest version number and an update timestamp.
19 In addition, link table 338 stores a deletion timestamp at delTime 348 to identify if
20 an account 342 is to be deleted from the link table.

21 A zero value for deletion timestamp 348 indicates that a value (i.e., an
22 account 342) is present in link table 338. A deletion timestamp 348 that indicates
23 a time identifies that the associated value 342 has been identified to be deleted
24 from the linked value list. That is, a non-zero value for deletion timestamp 348
25 indicates that a value is in an absent state and will not be rendered for display. A

1 deletion timestamp 348 is necessary as an identifier for record purposes when the
2 directory is replicated to indicate that a deletion of a value was performed at a
3 networked site. If the value is simply deleted and removed from the linked value
4 list without an identifier to indicate as such, there would be no record to update the
5 next directory when the network sites replicate.

6 **Multi-Valued Attribute Replication**

7 Fig. 5 illustrates how a replication conflict is avoided when two objects
8 having an updated multi-valued attribute are replicated in a network implementing
9 a linked value replication model. Initially, as shown in Fig. 4, domain controller A
10 has an object 314(A) with a multi-valued members attribute 334. The attribute has
11 two values, account1 and account2, in link table 338. Domain controller B also
12 has an up-to-date replica of object S.

13 In Fig. 5, a data administrator at domain controller A deletes account1 from
14 the security list 342(A) in link table 338(A). As illustrated, account1 is not
15 removed from link table 338(A), but rather identified as having been deleted.
16 Valuedata 346(A) for account1 is updated to version2 (v2) of the value occurring
17 at time2 (t2) as indicated by 500. To identify that account1 has been deleted,
18 deletion timestamp 348(A) is updated at time2 (t2) as indicated by 502.

19 The data administrator also adds a new account3 to the security list 342(A)
20 at domain controller A as indicated by 504. Valuedata 346(A) for account3 is
21 initialized to version1 (v1) of the value occurring at time3 (t3).

22 Within a replication latency period, a second data administrator at domain
23 controller B adds a new account 4 to the security list 342(B) as indicated by 506.
24 Valuedata 346(B) for account4 is initialized to version1 (v1) of the value
25 occurring at time4 (t4).

Fig. 6 illustrates that when domain controllers A and B replicate directories 312 and 320, respectively (Fig. 4), both of the value updates are accounted for in the resultant link table 338. Neither update is lost in resolving a replication conflict because the level of replication granularity is at the attribute value level, rather than at the attribute level. The update at domain controller A (delete account1 and add account3) and the update at domain controller B (add account4) do not cause a replication conflict because each account 342 has a different combination of version number and update timestamp in valuedata 346.

After domain controllers A and B replicate, and a designated period of time identified as the “tombstone lifetime”, the value account1 is removed (actually deleted) from link table 338 by a separate process that recognizes the value as having been identified for deletion. A tombstone lifetime is the period of time that deletions exist in a directory before being removed. The process of removing a value that has been identified for deletion is called “garbage collection”.

Link Collision

Figs. 7 and 8 illustrate that providing a creation timestamp for an attribute value 342 distinguishes incarnations of the values to avoid data loss during a “link collision”. A link collision occurs when a value is deleted (i.e., garbage collected) and then re-created within a replication latency period. A creation timestamp is included in valuedata 346 at the value level to prevent losing a re-created value during resolution of a replication conflict.

Initially, as shown in Fig. 4, domain controller A has an object 314(A) with a multi-valued members attribute 334. The attribute has two values, account1 and account2, in link table 338. Domain controller B also has an up-to-date replica of object S. Fig. 7 also shows domain controllers A and B each having an up-to-date

1 replica of object S. For simplification, only link table 338 for each object 314 is
2 shown in the figure.

3 A creation timestamp, identified with a “c”, is included in valuedata 346 for
4 each account 342 to indicate the creation time of each value. As shown, account1
5 was created at time c1 and version1 (v1) of account1 occurred at time1 (t1).
6 Account2 was created at time c2 and version3 (v3) of account2 occurred at time2
7 (t2). Creation timestamps can be derived independently without having to
8 correlate or synchronize time with other replicas stored on different computers.

9 Fig. 8 shows three instances of object 314(A) in domain controller A. At
10 instance 800, a data administrator at domain controller A deletes account2 from
11 the security list 342(A) in link table 338(A). Valuedata 346(A) for account 2 is
12 updated to version4 (v4) of the value occurring at time5 (t5) as indicated by 802.
13 To identify that account2 has been deleted, deletion timestamp 348(A) is updated
14 at time5 (t5) as indicated by 804.

15 At instance 806 of object 314(A) in domain controller A, the process of
16 garbage collection recognizes that account2 has been identified for deletion and
17 removes account2 from link table 338(A). The process of garbage collection
18 occurs before replication of domain controller A with domain controller B.

19 At instance 808 of object 314(A) in domain controller A, the data
20 administrator re-creates account2 which is added to the link table 342(A).
21 Valuedata 346(A) indicates that account2 was created at time c6 and version1 (v1)
22 of account2 occurred at time6 (t6). The version number is initialized as version1
23 because account2 is a new value added to the link table 338(A).

24 When domain controllers A and B replicate after account2 was deleted and
25 then re-created at domain controller A, there will be a replication conflict to

1 resolve because valuedata 348 for account2 has changed from the initial state of
2 c2, v3, t2 (Fig. 7) to c6, v1, t6 shown in Fig. 8 at 810. Without the creation
3 timestamp, the replica on domain controller B would prevail in the replication
4 conflict because account2 was initially identified as version3 (v3), and after
5 having been re-created, is identified as version1 (v1) on domain controller A. If
6 domain controller B prevails, the new account2 created at domain controller A
7 would be lost data. However, the replication conflict is resolved in favor of
8 domain controller A because creation timestamp c6 is later than the initial creation
9 timestamp c2 at domain controller B.

10 **Replication Transition from Attribute-Level to Attribute Value-Level**

11 Figs. 9, 10, and 11 illustrate an example of managing the architectural
12 incompatibilities between directory partitions that are replicated with a
13 conventional state-based replication model (i.e., replicated at the attribute level),
14 and updated directory partitions that can be replicated with the linked value
15 replication model described herein (i.e., replicated at the attribute value level).
16 The linked value replication model accounts for changes at both the attribute level
17 and the attribute value level to integrate the directory partitions for the two
18 replication models. Replication transition with the linked value replication model
19 does not require a manual database conversion, as is typically required of an
20 administrator when implementing a new database model. Conventional
21 replication at the attribute level is identified as “legacy replication” where
22 “legacy” defines a state-based replication model directory partition.

23 Fig. 9 shows a network 900 with three networked domain controllers 902,
24 904, and 906 (computers A, B, and C). The domain controllers are interconnected
25 via a communications network (not shown). The network 900 and domain

1 controllers A, B, and C are examples of the network 300 and domain controllers
2 302 described above and shown in Fig. 3.

3 The computers A, B, and C have a directory 908, 910, and 912,
4 respectively. Each directory stores a replica of a contact group 914 which contains
5 a group object 916. The group object 916, identified as object CG, is associated
6 with the contact group 914 and identifies individual clients in a contact list.

7 The group object 916 has attributes and metadata as described in relation to
8 object 314 shown in Fig. 4. The object 916 has a multi-valued members attribute
9 918 that associates the individual clients in the contact list. Metadata 920 for the
10 members attribute includes a link table reference to a data structure 922. Link
11 table 922 maintains the linked values (e.g., the clients 924 in the contact list) for
12 the multi-valued members attribute 918.

13 Link table 922 maintains valuedata 926 and a deletion timestamp 928 for
14 each client 924. The valuedata 926, delTime 928, and other aspects of link table
15 922 are also described in relation to link table 338 shown in Fig. 4.

16 Computers A, B, and C initially have a legacy directory replica of object
17 916 that has a multi-valued members attribute 918 which has two values, client1
18 and client2. In an initial legacy mode, metadata 920 includes a latest version
19 number, v1, and an update timestamp, t1, for the members attribute 918. Also for
20 an initial legacy mode, valuedata 926 for each value (i.e., the clients 924) is null,
21 or zero, and the deletion timestamp 928 is zero to indicate the existence of a
22 particular value.

23 Fig. 10 shows an instance of object 916 in each of the computers A, B, and
24 C. For simplification, only the link table 922, members attribute 918, and
25 metadata 920 for the members attribute is shown in the figure for each object 916.

1 In this example, computers A and B implement the linked value replication model
2 (i.e., “new mode”) described above with respect to Figs. 4, 5, and 6. Computer C
3 implements the conventional state-based replication model (i.e., “legacy mode”).

4 At computer A, a data administrator adds a new client3 in link table
5 922(A). Because computer A implements linked value replication, valuedata
6 926(A) for client3 is initialized to version1 (v1) of the value occurring at time2
7 (t2). For a linked value replication model, non-null valuedata is a non-zero value
8 (i.e., valuedata 926(A) for client 3). That is, a version of a linked value is one or
9 more and valid timestamp is non-zero. Existent, or non-null, valuedata
10 distinguishes a linked value replication model over an attribute replication model.
11 In the case of a replication conflict, a linked value having non-null valuedata will
12 prevail over a linked value having null valuedata. This establishes a resolution
13 policy that values having conflict resolution data prevail over values without
14 conflict resolution data.

15 At computer B, a data administrator deletes client2 from link table 922(B).
16 Because computer B implements linked value replication, the deletion timestamp
17 928(B) for client2 is updated to time3 (t3) to indicate that the value has been
18 identified for deletion. Valuedata 926(B) updates from the null value to version1
19 (v1) of the value occurring at time3 (t3).

20 At computer C, a data administrator deletes client1 from link table 922(C).
21 Because computer C is operating in the legacy mode of state-based replication,
22 client1 is actually removed from link table 922(C), rather than being identified for
23 deletion at the value level with a deletion timestamp. In the legacy mode of state-
24 based replication, the value level data is not created. Rather, the attribute level
25

1 metadata 920(C) is updated to version2 (v2) of the attribute occurring at time4 (t4)
2 to indicate that a value of the members attribute 918(C) has been changed.

3 Fig. 11 shows the results of computers A, B, and C replicating after the
4 changes to the values in link tables 922(A), 922(B), and 922(C), respectively.
5 Domain controllers (servers, computers, etc.) operating with the linked value
6 replication model cannot replicate from domain controllers operating under the
7 legacy mode of state-based replication. That is, computers A and B cannot
8 replicate from computer C. However, computer C can replicate from computers A
9 and B, but has to first “promote” itself to the new mode prior to replicating with
10 either computer A or B. Computer C promotes itself to implement linked value
11 replication when it first replicates with a computer in the network operating with
12 the linked value replication model.

13 Replication transition from attribute level to attribute value level occurs in
14 two stages: first at the attribute level (i.e., conventional “legacy” replication), and
15 second at the attribute value level. At the attribute level, attributes having a later
16 version number and/or timestamp are replicated first. This stage of the replication
17 includes only those linked values that do not have valuedata. Subsequently, at the
18 value level, values having more recent valuedata are replicated second. With
19 replication transition, values having null valuedata are included in the attribute
20 level replication stage and excluded from the value level replication stage.

21 In Fig. 11, computer C first replicates with computer B. Client2 exists on
22 computer C as a legacy value (i.e., valuedata 926(C) and delTime 928(C) for
23 client2 is null, Fig. 10). When replicating with computer B, computer B prevails
24 in a replication conflict because client2 has value level data. Computer C updates

1 valuedata 926(C) and delTime 928(C) for client2 to indicate that the value has
2 been identified to be deleted.

3 Computer C next replicates with computer A and adds client3 to link table
4 922(C). Valuedata 926(C) is initialized to version1 (v1) of client3 occurring at
5 time2 (t2). Computer C does not replicate client1 from computer A because
6 client1 is a legacy value having no value level data.

7 Computer B replicates from computer C and updates the change to the
8 members attribute metadata 920(B) to reflect the update made in computer C.
9 Computer B then accounts for updates and changes at the attribute level (i.e.,
10 members attribute 918(B)), and replicates only legacy values without any value
11 level data from computer C. This follows the conventional state based replication
12 model. However, computer C does not have any legacy values without value level
13 data, but rather has client2 and client3 each with valuedata 926(C). Thus,
14 computer B receives an empty list from computer C with no legacy value changes
15 to be made. This indicates to computer B to remove any local legacy values from
16 the link table. Accordingly, computer B removes client1 from link table 922(B).

17 After accounting for attribute level replication, computer B replicates at the
18 value level implementing the link value replication model. Computer B adds
19 client3 from computer C to link table 922(B) and initializes valuedata 926(B).
20 Computer B does not replicate from computer A because computer B is
21 transitively updated from computer A. Computer C replicates from computer A
22 before computer B replicates from computer C.

23 Computer A replicates from computer B and updates the change to
24 members attribute metadata 920(A) to reflect the update made in computer B,
25 which was initiated in computer C. Computer A then accounts for updates and

1 changes at the attribute level (i.e., members attribute 918(A)), and replicates only
2 legacy values without any value level data from computer B. However, computer
3 B does not have any legacy values without value level data, but rather has client2
4 and client3 each with valuedata 926(B). Thus, computer A receives an empty list
5 from computer B with no legacy value changes to be made. This indicates to
6 computer A to remove any local legacy values. Accordingly, computer A removes
7 client1 and client 2 from link table 922(A).

8 After accounting for attribute level replication, computer A replicates at the
9 value level implementing the link value replication model. Computer A adds
10 client2 (which does not exist because it was just removed) from computer B to
11 link table 922(A) and updates valuedata 926(A) and delTime 928(A) to indicate
12 that client2 has been identified to be deleted. Computer A does not replicate from
13 computer C because computer A is transitively updated from computer C.
14 Computer B replicates from computer C before computer A replicates from
15 computer B.

16 Fig. 11 shows that computers A, B, and C, have all converged to the same
17 set of values via the link value replication model. The example illustrates how
18 directory partitions are replicated from an existing attribute level to a linked value
19 level. The link value replication model reduces the amount of data that is
20 communicated between domain controllers in a network when replicating
21 directory partitions, reduces the possibilities of replication convergence conflicts,
22 and provides architectural compatibility between a conventional state-based
23 replication model and the link value replication model.

1 Fig. 12 illustrates a method to replicate multi-valued object attributes
2 having attribute-value level conflict-resolution data. At block 400, an object
3 stored in a first directory at a network domain controller is replicated with a
4 replica of the object stored in a second directory at a second network domain
5 controller. The object has a multi-valued attribute comprised of individual values
6 each having associated conflict-resolution data.

7 At block 402, the conflict-resolution data for the individual values of the
8 object stored in the first directory and of the replica of the object stored in the
9 second directory is compared to determine if a replication conflict exists between
10 the individual values. At block 404, a creation timestamp for the individual values
11 is compared to determine if an attribute value, or the replica of the attribute value,
12 has changed.

13 If the creation timestamp indicates that one of the values was created after
14 the other (i.e., “yes” from block 404), the attribute value having the earlier
15 creation timestamp is updated with the attribute value that has the later creation
16 timestamp at block 406. That is, the older value created first is replicated with any
17 associated data from the newer value that was created last. If the creation
18 timestamp is the same for the two values (i.e. “no” from block 404), a version
19 number for the individual values is compared to determine if an attribute value, or
20 the replica of the attribute value, has been updated or changed to a new version at
21 block 408.

22 If the version number indicates that one of the values was updated or
23 changed to a more recent version (i.e., “yes” from block 408), the attribute value
24 having the lower version number is updated with the attribute value that has the
25

higher version number at block 410. That is, the older value with the lower version number is replicated with any associated data from the newer value that was updated or changed last. If the version number is the same for the two values (i.e., “no” from block 408), an update timestamp for the individual values is compared to determine if an attribute value, or the replica of the attribute value, has been updated at block 412.

If the update timestamp indicates that one of the values was updated or changed after the other (yet the version number remains the same) (i.e., “yes” from block 412), the attribute value having the earlier update timestamp is updated with the attribute value that has the later update timestamp at block 414. That is, the older value is replicated with any associated data from the newer value that was updated or changed last. If the update timestamp is the same for the two values (i.e. “no” from block 412), then there is no replication conflict to be resolved between the individual values of the multi-valued object attribute (block 416).

At block 418, a deletion timestamp is evaluated to determine if an individual value has been identified to be deleted. If the deletion timestamp is not null (i.e., “no” from block 418), then the value is deleted from the object attribute at block 420. That is, if a value has been identified to be deleted from the object attribute, then the deletion timestamp will indicate when the value was marked for deletion. If the deletion timestamp indicates null (i.e., “yes” from block 418), then the method continues to replicate directory objects (at block 400).

Exemplary Computing System and Environment

Fig. 13 illustrates an example of a computing environment 500 within which the computer, network, and system architectures described herein can be

either fully or partially implemented. Exemplary computing environment 500 is only one example of a computing system and is not intended to suggest any limitation as to the scope of use or functionality of the network architectures. Neither should the computing environment 500 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary computing environment 500.

The computer and network architectures can be implemented with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use include, but are not limited to, personal computers, server computers, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

Link value replication may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Link value replication may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

1 The computing environment 500 includes a general-purpose computing
2 system in the form of a computer 502. The components of computer 502 can
3 include, by are not limited to, one or more processors or processing units 504, a
4 system memory 506, and a system bus 508 that couples various system
5 components including the processor 504 to the system memory 506.

6 The system bus 508 represents one or more of any of several types of bus
7 structures, including a memory bus or memory controller, a peripheral bus, an
8 accelerated graphics port, and a processor or local bus using any of a variety of
9 bus architectures. By way of example, such architectures can include an Industry
10 Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an
11 Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA)
12 local bus, and a Peripheral Component Interconnects (PCI) bus also known as a
13 Mezzanine bus.

14 Computer system 502 typically includes a variety of computer readable
15 media. Such media can be any available media that is accessible by computer 502
16 and includes both volatile and non-volatile media, removable and non-removable
17 media. The system memory 506 includes computer readable media in the form of
18 volatile memory, such as random access memory (RAM) 510, and/or non-volatile
19 memory, such as read only memory (ROM) 512. A basic input/output system
20 (BIOS) 514, containing the basic routines that help to transfer information
21 between elements within computer 502, such as during start-up, is stored in ROM
22 512. RAM 510 typically contains data and/or program modules that are
23 immediately accessible to and/or presently operated on by the processing unit 504.

24 Computer 502 can also include other removable/non-removable,
25 volatile/non-volatile computer storage media. By way of example, Fig. 13

1 illustrates a hard disk drive 516 for reading from and writing to a non-removable,
2 non-volatile magnetic media (not shown), a magnetic disk drive 518 for reading
3 from and writing to a removable, non-volatile magnetic disk 520 (e.g., a “floppy
4 disk”), and an optical disk drive 522 for reading from and/or writing to a
5 removable, non-volatile optical disk 524 such as a CD-ROM, DVD-ROM, or other
6 optical media. The hard disk drive 516, magnetic disk drive 518, and optical disk
7 drive 522 are each connected to the system bus 508 by one or more data media
8 interfaces 526. Alternatively, the hard disk drive 516, magnetic disk drive 518,
9 and optical disk drive 522 can be connected to the system bus 508 by a SCSI
10 interface (not shown).

11 The disk drives and their associated computer-readable media provide non-
12 volatile storage of computer readable instructions, data structures, program
13 modules, and other data for computer 502. Although the example illustrates a
14 hard disk 516, a removable magnetic disk 520, and a removable optical disk 524,
15 it is to be appreciated that other types of computer readable media which can store
16 data that is accessible by a computer, such as magnetic cassettes or other magnetic
17 storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or
18 other optical storage, random access memories (RAM), read only memories
19 (ROM), electrically erasable programmable read-only memory (EEPROM), and
20 the like, can also be utilized to implement the exemplary computing system and
21 environment.

22 Any number of program modules can be stored on the hard disk 516,
23 magnetic disk 520, optical disk 524, ROM 512, and/or RAM 510, including by
24 way of example, an operating system 526, one or more application programs 528,
25 other program modules 530, and program data 532. Each of such operating

1 system 526, one or more application programs 528, other program modules 530,
2 and program data 532 (or some combination thereof) may include an embodiment
3 of link value replication.

4 Computer system 502 can include a variety of computer readable media
5 identified as communication media. Communication media typically embodies
6 computer readable instructions, data structures, program modules, or other data in
7 a modulated data signal such as a carrier wave or other transport mechanism and
8 includes any information delivery media. The term “modulated data signal”
9 means a signal that has one or more of its characteristics set or changed in such a
10 manner as to encode information in the signal. By way of example, and not
11 limitation, communication media includes wired media such as a wired network or
12 direct-wired connection, and wireless media such as acoustic, RF, infrared, and
13 other wireless media. Combinations of any of the above are also included within
14 the scope of computer readable media.

15 A user can enter commands and information into computer system 502 via
16 input devices such as a keyboard 534 and a pointing device 536 (e.g., a “mouse”).
17 Other input devices 538 (not shown specifically) may include a microphone,
18 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and
19 other input devices are connected to the processing unit 604 via input/output
20 interfaces 540 that are coupled to the system bus 508, but may be connected by
21 other interface and bus structures, such as a parallel port, game port, or a universal
22 serial bus (USB).

23 A monitor 542 or other type of display device can also be connected to the
24 system bus 508 via an interface, such as a video adapter 544. In addition to the
25 monitor 542, other output peripheral devices can include components such as

1 speakers (not shown) and a printer 546 which can be connected to computer 502
2 via the input/output interfaces 540.

3 Computer 502 can operate in a networked environment using logical
4 connections to one or more remote computers, such as a remote computing device
5 548. By way of example, the remote computing device 548 can be a personal
6 computer, portable computer, a server, a router, a network computer, a peer device
7 or other common network node, and the like. The remote computing device 548 is
8 illustrated as a portable computer that can include many or all of the elements and
9 features described herein relative to computer system 502.

10 Logical connections between computer 502 and the remote computer 548
11 are depicted as a local area network (LAN) 550 and a general wide area network
12 (WAN) 552. Such networking environments are commonplace in offices,
13 enterprise-wide computer networks, intranets, and the Internet. When
14 implemented in a LAN networking environment, the computer 502 is connected to
15 a local network 550 via a network interface or adapter 554. When implemented in
16 a WAN networking environment, the computer 502 typically includes a modem
17 556 or other means for establishing communications over the wide network 552.
18 The modem 556, which can be internal or external to computer 502, can be
19 connected to the system bus 508 via the input/output interfaces 540 or other
20 appropriate mechanisms. It is to be appreciated that the illustrated network
21 connections are exemplary and that other means of establishing communication
22 link(s) between the computers 502 and 548 can be employed.

23 In a networked environment, such as that illustrated with computing
24 environment 500, program modules depicted relative to the computer 502, or
25 portions thereof, may be stored in a remote memory storage device. By way of

1 example, remote application programs 558 reside on a memory device of remote
2 computer 548. For purposes of illustration, application programs and other
3 executable program components, such as the operating system, are illustrated
4 herein as discrete blocks, although it is recognized that such programs and
5 components reside at various times in different storage components of the
6 computer system 502, and are executed by the data processor(s) of the computer.

7 **Conclusion**

8 Although the systems and methods have been described in language
9 specific to structural features and/or methodological steps, it is to be understood
10 that the technology defined in the appended claims is not necessarily limited to the
11 specific features or steps described. Rather, the specific features and steps are
12 disclosed as preferred forms of implementing the claimed invention.

TOP SECRET//COMINT